

CLAIMS:

1. A document search system, comprising:
 - a document index comprising a plurality of document index partitions, each partition comprising a subset of the document index; each document index partition mapping terms to documents;
 - a partition index that maps a specified term to a subset of the document index partitions, wherein the document index partitions in the subset each map the specified term to at least one document containing the specified term;
 - a balancer configured to receive a search query having a set of terms, comprising one or more terms, to search the partition index so as to identify a subset of the document index partitions that potentially include documents that satisfy the search query, and to direct the search query to only the identified subset of the document index partitions.
2. The document search system of claim 1, wherein the partition index maps the specified term to an empty subset when the specified term is not contained within the partition index.
3. The document search system of claim 1, including one or more servers configured to search the document index partitions in the subset for terms in the search query.
4. The document search system of claim 1, including a plurality of index search servers, each index search server configured to search at least a portion of at least one document index partition of the plurality of document index partitions so as to identify documents containing specified terms.
5. The document search system of claim 4, wherein the balancer is configured to direct the search query to only the servers, of the one of more servers, that are configured to search document index partitions included in the identified subset.
6. The document search system of claim 1, wherein the balancer is configured so that when the search query includes a plurality of distinct terms, the balancer searches the partition index so as to obtain a plurality of maps, performs a Boolean operation on the plurality of maps so as to generate a single map, and based on the single map identifies the subset of document index partitions.

7. The document search system of claim 1, wherein the balancer comprises:
 - a processor;
 - a communications interface; and
 - a memory, comprising:
 - communications procedures for receiving the search query, and for transmitting search results; and
 - a balancer filter comprising:
 - mapping instructions for searching the partition index for each term in the search query so as to generate a map for each said term;
 - combining instructions, utilized when the search query comprises a plurality of terms and the first instructions generate a plurality of maps, for generating a single map from the plurality of maps; and
 - identifying instructions for identifying a subset of the document index partitions in accordance with the map or single map; and
 - distribution instructions for sending the search query to each document index partition, if any, in the identified subset.
8. The document search system of claim 1, wherein
 - the partition index is configured to map the specified term to a set of document index sub-partitions, each document index sub-partition corresponding to a subset of the documents indexed by a respective document index partition of the plurality of document index partitions; wherein each document index sub-partition to which the specified term is mapped by the partition index maps the specified term to at least one document having the specified term.
9. The document search system of claim 8, wherein
 - the balancer is configured so that when the search query includes a plurality of distinct terms, the balancer searches the partition index so as to obtain a plurality of maps, performs a Boolean operation on the plurality of maps so as to generate a single map, wherein the single map identifies a set of document index sub-partitions, and based on the identified set of document index sub-partitions identifies the subset of document index partitions.

10. The document search system of 8, wherein
 - a plurality of document index sub-partitions correspond to each document index partition of the plurality of document index partitions; and
 - the balancer is configured so to search the partition index so as to obtain a map for each term of the search query, the map identifying a subset of the document index sub-partitions, wherein each document index sub-partition in the subset contains at least one document having the specified term.
11. The document search system of claim 10, including a plurality of index search servers, each index search server configured to search at least one document index sub-partition so as to identify documents containing specified terms.
12. The document search system of claim 11, wherein the map identifying the subset comprises a set of bits, each respective bit of the map corresponds to a respective subset of the index search servers, and the balancer is configured to direct the search query to only index search servers corresponding to bits in the map having a first predefined value.
13. The document search system of claim 11, wherein the balancer is configured to direct the search query to only the index search servers, of the plurality of index search servers, that are configured to search document index sub-partitions included in the identified subset of document index sub-partitions.
14. A method of searching for documents, comprising:
 - receiving a search query containing a set of search terms, the set of search terms having at least one search term;
 - searching for the set of search terms in a partition index so as to identify a subset of a plurality of document index partitions; and
 - searching, in only those document index partitions in the identified subset, for documents containing the set of search terms;
 - wherein the partition index maps any specified term to a respective subset of the document index partitions, wherein each document index partition in the respective subset maps the specified term to at least one document containing the specified term.

15. The method of claim 14, wherein the searching in the document index partitions in the identified subset includes directing the search query to one or more servers configured to search the document index partitions in the identified subset.

16. The method of claim 14, wherein the searching in the partition index includes, when the search query includes a plurality of distinct terms, searching the partition index with respect to each term in the plurality of distinct terms so as to obtain a plurality of maps, performing a Boolean operation on the plurality of maps so as to generate a single map, and based on the single map identifying the subset of document index partitions.

17. The method of claim 14, wherein the partition index is configured to map a specified term in the search query to a set of document index sub-partitions, each document index sub-partition corresponding to a subset of the documents indexed by a respective document index partition of the set of document index partitions; wherein each document index sub-partition to which the specified term is mapped by the partition index maps the specified term to at least one document having the specified term.

18. The method of claim 17, wherein the searching in the partition index includes, when the search query includes a plurality of distinct terms, searching the partition index with respect to each term in the plurality of distinct terms so as to obtain a plurality of maps, performing a Boolean operation on the plurality of maps so as to generate a single map, and based on the single map identifying the subset of document index partitions, wherein the single map identifies a set of document index sub-partitions.

19. The method of claim 17, including searching, in only those document index sub-partitions in the identified set of document index sub-partitions, for documents containing the set of search terms.

20. The method of claim 17, including identifying the subset of document index partitions based on the identified set of document index sub-partitions.

21. The method of claim 17, wherein
a plurality of document index sub-partitions correspond to each document index partition of the plurality of document index partitions; and
the method includes searching the partition index so as to obtain a map for each term of the search query, the map identifying a subset of the document index

- sub-partitions, wherein each document index sub-partition in the subset maps the specified term to at least one document having the specified term.
22. The method of claim 21, including searching, in only those document index sub-partitions in the identified subset of document index sub-partitions, for documents containing the set of search terms.
23. The method of claim 21, wherein the map identifying the subset comprises a set of bits, each respective bit of the map corresponds to a respective subset of a set of index search servers, and the searching for documents includes directing the search query only to index search servers corresponding to bits in the map having a first predefined value.
24. The method of claim 21, including directing the search query only to index search servers, of a plurality of index search servers, that are configured to search document index sub-partitions included in the identified subset of document index sub-partitions.
25. A document search system, comprising:
- a document index comprising a plurality of document index partitions, each partition comprising a subset of the document index; each respective document index partition comprising a plurality of document index sub-partitions, each document index sub-partition comprising a subset of the respective document index partition; each document index sub-partition configured to map terms to documents;
 - a plurality of partition indexes, each respective partition configured to maps a specified term to a subset of the document index sub-partitions of a corresponding document index partition, wherein the document index sub-partitions in the subset each map the specified term to at least one document containing the specified term;
 - a plurality of balancers, each associated with a respective partition index of the plurality of partition indexes; each respective balancer configured to receive a search query having a set of terms, comprising one or more terms, to search the partition index associated with the respective balancer so as to identify a subset of the document index sub-partitions that potentially include documents that satisfy the search query, and to direct the search query to only the identified subset of the document index sub-partitions.
26. A method of searching for documents, comprising:

receiving a search query containing a set of search terms, the set of search terms having at least one search term;

searching for the set of search terms in a partition index so as to identify a subset of a plurality of document index sub-partitions; and

searching, in only those document index sub-partitions in the identified subset, for documents containing the set of search terms;

wherein the partition index maps any specified term to a respective subset of the document index sub-partitions, wherein each document index sub-partition in the respective subset maps the specified term to at least one document containing the specified term.

27. A method of searching for documents, comprising:

receiving a search query containing a set of search terms, the set of search terms having at least one search term;

searching for the set of search terms in each of a plurality of partition indexes, each partition index corresponding to a document index partition of a document index, wherein the search of each respective partition index identifies a subset of a plurality of document index sub-partitions of the document index partition corresponding to the respective partition index; and

searching, in only those document index sub-partitions in the identified subsets, for documents containing the set of search terms;

wherein each respective partition index maps any specified term to a respective subset of the document index sub-partitions of a corresponding document index partition, wherein each document index sub-partition in the respective subset maps the specified term to at least one document containing the specified term.